When Performance Portability is less than Perfect

Matching Applications to Architectures

Mark O'Connor

Director Product Management, HPC & Server Tools

CoE PP, Denver
08 / 2017

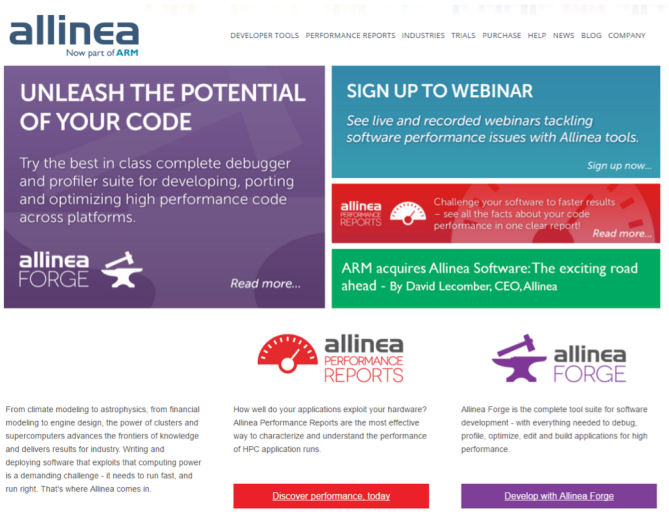

Allinea: leaders in cross-platform developer tools for HPC

ARM

Portable performance ≠ equivalent performance

Case study: Haswell vs KNL

Different applications benefit from the KNL architecture to different degrees

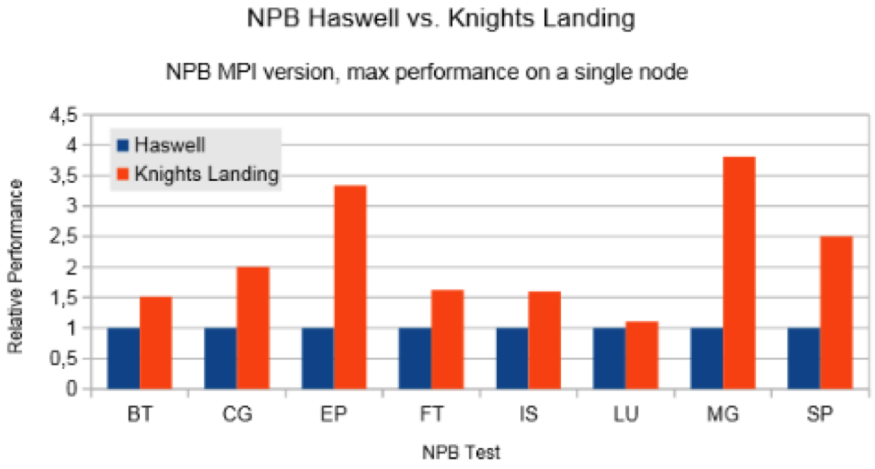

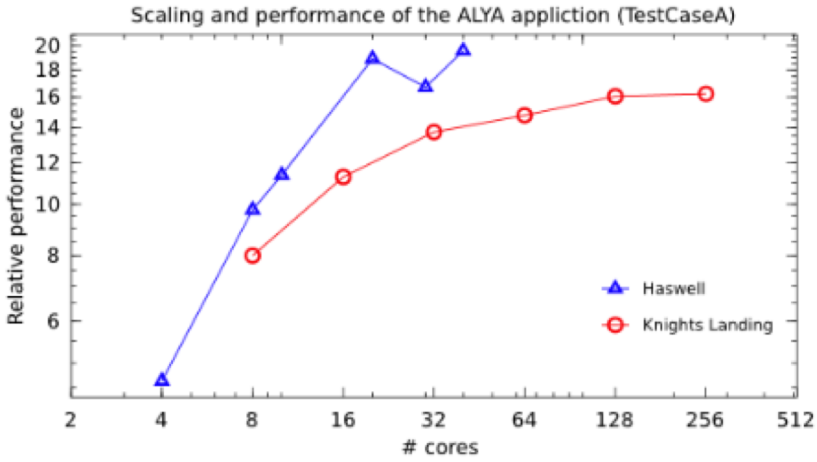

Source: http://www.prace-ri.eu/best-practice-guide-knights-landing-january-2017/

Can we predict application performance?

Cori saw I/O performance differences when changing CPU architecture

Microkernels work for single-core optimization, not application characterization

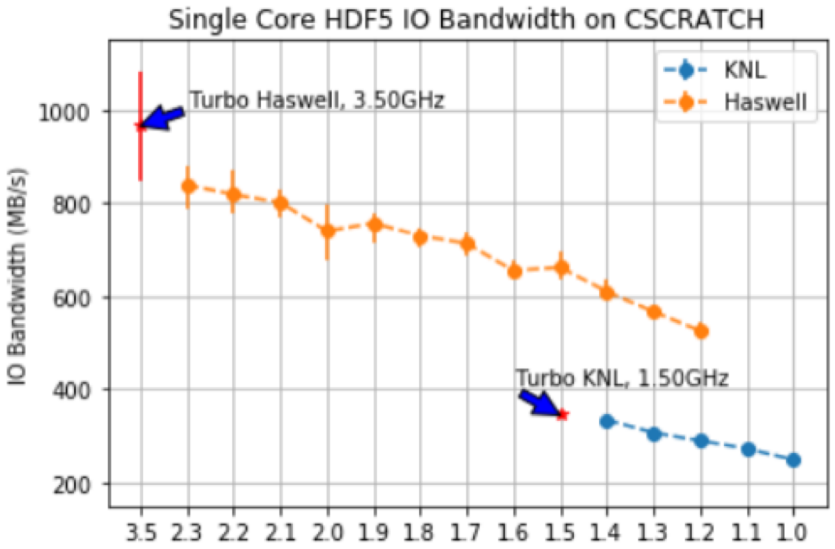

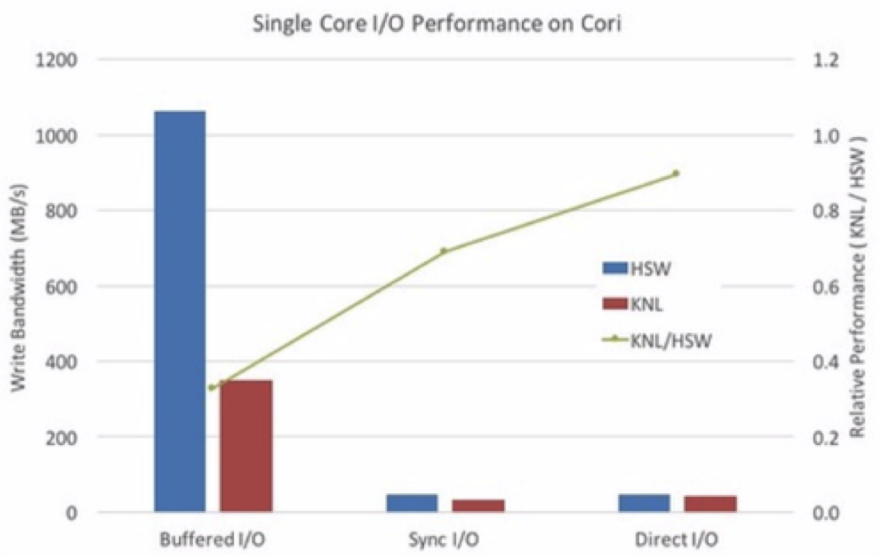

Source:
http://www.francoistessier.info/documents/CUG2017.pdf

Current attempts at application characterization

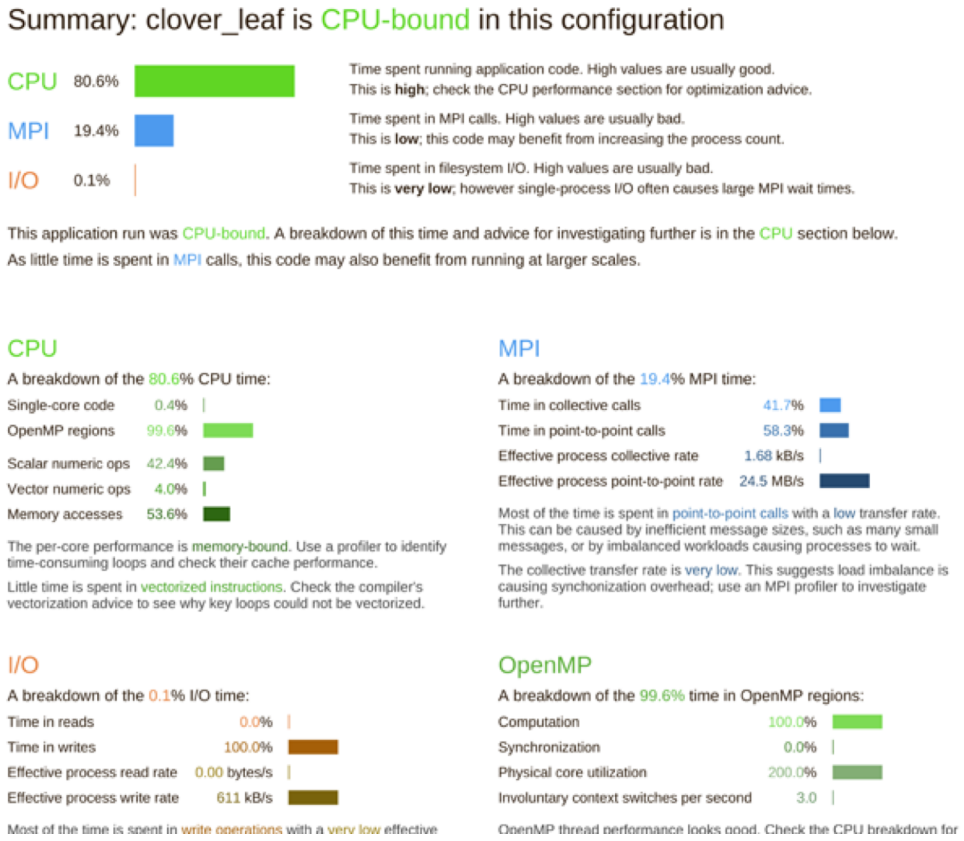

Single page summary per application

- CPU vs MPI vs I/O breakdown
- Time in vector and memory instructions
- Effective MPI and I/O bandwidth
- Memory usage
- OpenMP / threading efficiency

Also available in Allinea MAP, plus:

- Breakdown over time and per source line:

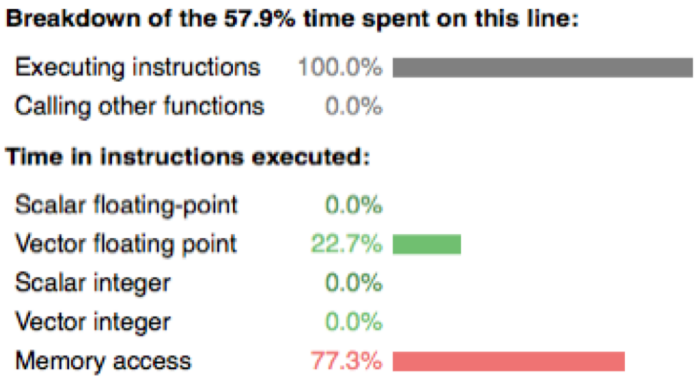

ARM

Standardizing application characterization

Characterizing CPU performance – the roofline model

FLOPs/second plotted against FLOPs/byte – very successful for compute characteristics

A CPU-centric view of application performance. Can we extend it for storage hierarchies?

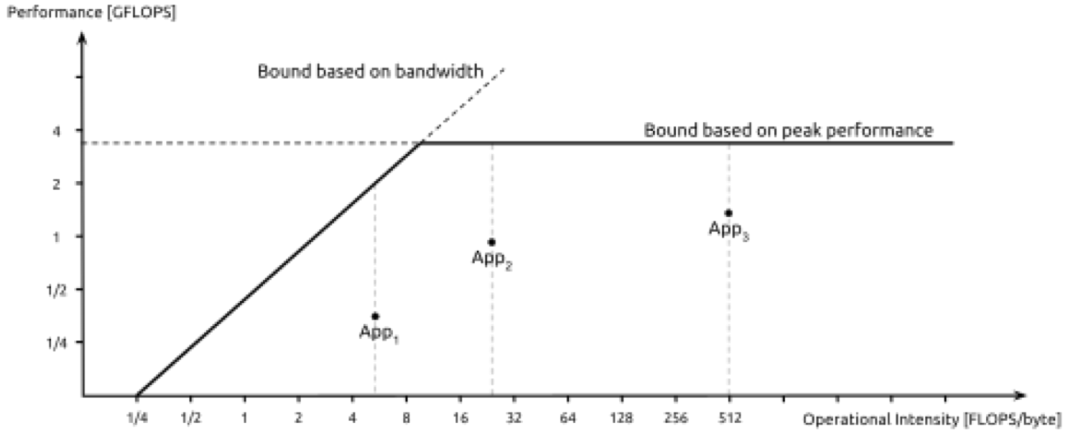

What would the roofline model for application portability look like?

Can we boil it down to two dimensions? Or a series of plots looking at different aspects?

ARM

Cache-aware roofline model

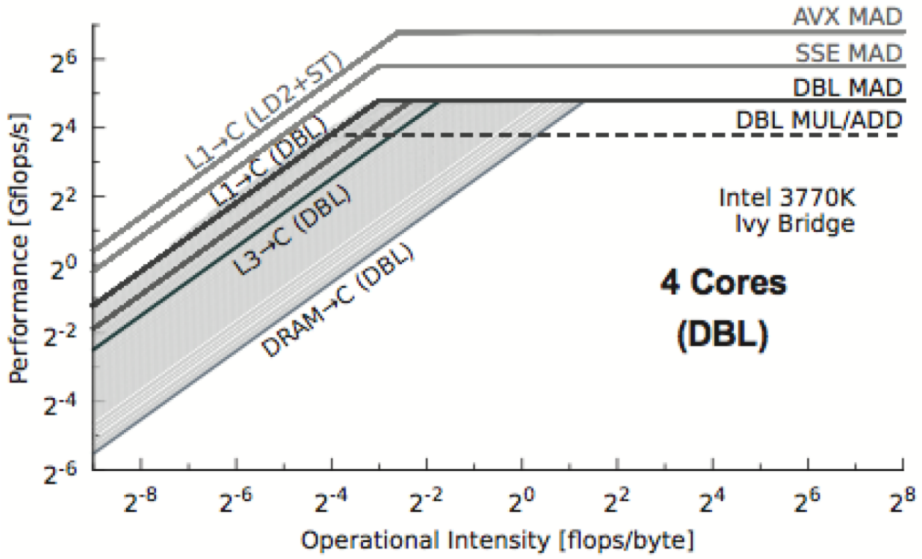

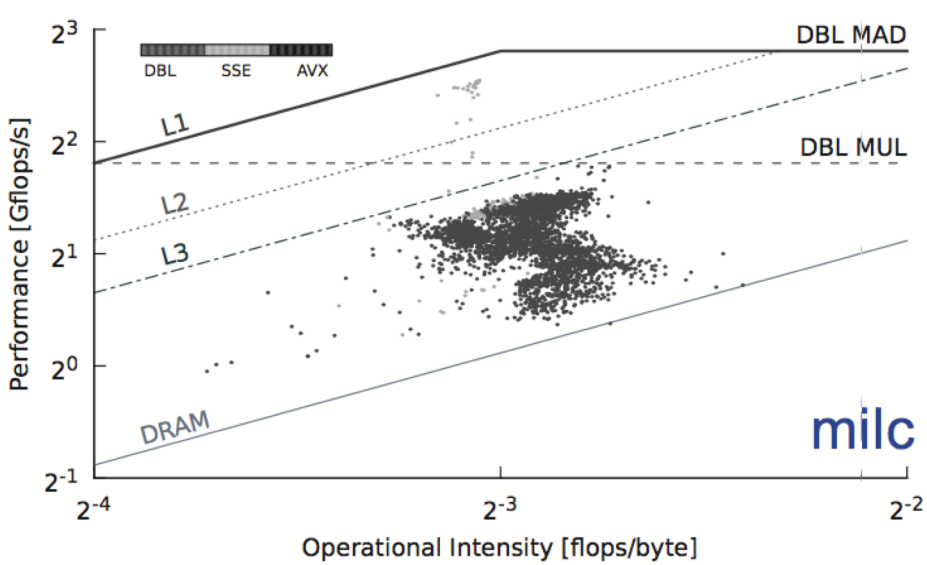

Extend to I/O-aware roofline model? Accelerator-aware roofline model?

If we want this, we need specific, reliable cross-manufacture performance counters.

Moving away from using as a loop-optimization tool to characterizing many applications on a cluster.

Source: http://sips.inesc-id.pt/~ilic/roofline.php

Exploring further extensions to the model

What happens if we tweak the vertical axis?

Performance isn't just Gflops/s. Some interesting alternatives:

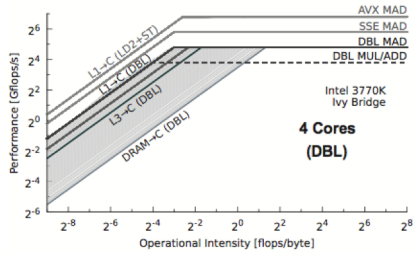

Gflops/cycle

Does this allow us to compare different architectures more meaningfully?

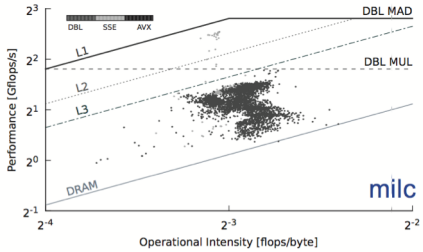

Gflops/joule

What if energy to solution across the cluster is more important than time to solution for each individual run?

Gflops/$

Can we now compare cloud offerings and bare-metal runs? Do our horizontal bars now show the trade-off in bursting to a cloud provider?

ARM

Other options to match applications to hardware

Machine learning

- Train an autoencoder to reduce dimensionality to 2D or 3D
- Cluster and predict

Reinforcement learning

- Train an end-to-end network to optimally place applications on clusters
- Plenty of training data available

Training and best practices

- Let *users* decide where to run their code
- Training humans may be harder than training machines

Thank-you!
Questions and discussion